

Multiobjective Optimization and Phase Transitions

Luís F Seoane^{1,2} and Ricard Solé^{1,2,3}

¹ ICREA-Complex Systems Lab, Universitat Pompeu Fabra – PRBB, Dr. Aiguader 88, 08003 Barcelona, Spain.

² Institut de Biologia Evolutiva, CSIC-UPF, Pg. Maritim de la Barceloneta 37, 08003 Barcelona.

³ Santa Fe Institute, 1399 Hyde Park Road, Santa Fe NM 87501, USA.
{luis.seoane,ricard.sole}@upf.edu

Abstract. Many complex systems obey to optimality conditions that are usually not simple. Conflicting traits often interact making a Multi Objective Optimization (MOO) approach necessary. Recent MOO research on complex systems report about the Pareto front (optimal designs implementing the best trade-off) in a qualitative manner. Meanwhile, research on traditional Simple Objective Optimization (SOO) often finds phase transitions and critical points. We summarize a robust framework that accounts for phase transitions located through SOO techniques and indicates what MOO features resolutely lead to phase transitions. These appear determined by the shape of the Pareto front, which at the same time is deeply related to the thermodynamic Gibbs surface. Indeed, thermodynamics can be written as an MOO from where its phase transitions can be parsimoniously derived; suggesting that the similarities between transitions in MOO-SOO and Statistical Mechanics go beyond mere coincidence.

Keywords: multiobjective optimization, Pareto optimality, phase transitions, statistical mechanics, thermodynamics

1 Introduction

Optimization has always been a major topic in complex systems research. Optimality conditions are relevant for a wealth of biological [1], [2], [3], [4], [5] and other natural and synthetic systems [6], [7], [8], [9], [10]. Evolution through natural selection is a main driver of biological systems towards optimal designs [11] [12] and certain physical principles (e.g. maximum entropy or optimal diffusion structures) already introduce a bias towards functional extrema. Human-made systems are equally constrained through cost-efficiency calculations – e.g. in transportation networks [9], [10].

Describing these situations requires optimal designs that often cope with interacting constraints. To give a good account of these selective forces, a *Pareto* or *Multi Objective Optimization* (MOO) approach can be useful. Let us introduce

this theory through a recent relevant example [13]. (Technical definitions follow below.) Consider the set (Γ) of all connected networks with a fixed number of nodes ($\gamma \in \Gamma$, figure 1a). Among them we seek those minimizing the average path length $\langle l \rangle(\gamma)$ and the number of edges $\rho(\gamma)$. These are the *target functions* ($T_f(\gamma) \equiv \{t_1 = \langle l \rangle(\gamma), t_2 = \rho(\gamma)\}$) of our MOO problem. A fully connected clique minimizes the average path length, but we need to implement all possible links, which is costly. The minimum spanning tree has the least number of edges possible but its average path length is quite large. Take networks γ_1 and γ_2 trading between these extremes and such that $\langle l \rangle(\gamma_1) < \langle l \rangle(\gamma_2)$ and $\rho(\gamma_1) < \rho(\gamma_2)$. This means that γ_1 implements a better tradeoff than γ_2 and we say then that γ_1 dominates γ_2 (figure 1c). A network ($\gamma_\pi \in \Pi$) $\subset \Gamma$ not dominated by any other $\gamma \in \Gamma$ is *Pareto optimal*. Often we cannot choose between a pair of networks because one is better than the other with respect to a target and worst with respect to the other – i.e. they are mutually not dominated. Because this situation is common, MOO solutions are often not a single global optimizer, but the collection of Pareto optimal (mutually non-dominated) networks that implement the most optimal tradeoff possible. We name this *Pareto optimal set* $\Pi \subset \Gamma$.

Target functions map each network $\gamma \in \Gamma$ into a point of \mathbb{R}^2 : $(\langle l \rangle(\gamma), \rho(\gamma))$. This plane, with the relevant traits in its axes, constitutes a *morphospace* in which salient network topologies are located as a function of their morphology [14] (figure 1b). Morphospace of other systems visualize phenotypes or designs with respect to relevant properties. The Pareto optimal set is mapped onto $T_f(\Pi)$ and constitutes a boundary of the morphospace (figure 1b-c). also known as the Pareto front.

Some authors are beginning to explore the consequences of Pareto optimality in biological systems [3], [4], [5] or in relevant models such as networks [15], [16] or regulatory circuits [5], [17]. While they tackle relevant questions through MOO methods, the description of these optimal designs is often a qualitative account of the elements along the Pareto front (as in the study of a restricted morphospace – an interesting contribution nevertheless). The same qualitative bias appears in classic MOO literature. Is a more quantitative analysis possible? Are there *universal features* that reach through different MOO problems, thus uniting Pareto optimal systems despite their differences? Through our research [18] (sketched in section 2) we have found a connection between MOO and statistical mechanics. Those *universal features* we were looking for are phase transitions and critical points, which leave clear imprints in the shape of the Pareto front. Some authors had explored MOO with Single Objective Optimization (SOO) methods – e.g. by integrating all targets linearly to define a *global energy function* $\Omega(\Lambda) = \sum_k \lambda_k t_k$, with $\Lambda = \{\lambda_1, \dots, \lambda_k\}$ arbitrary parameters that introduce a bias towards some of the targets. Such research often finds phase transitions and other phenomena [8], [10]. A parsimonious theory lacked as to why some systems would present such transitions and others would not.

To the best of our knowledge, authors researching MOO do not exploit this connection with thermodynamics which, we believe, much enriches the discus-

sion of Pareto optimal designs. Two relevant examples from network theory: The efficiency of different topologies has been researched for the relay of information across a network using two distinct delivery heuristics [15]. It was made an exhaustive work in describing network topologies and locating them in a morphospace in which different network features are segregated. In that same morphospace, the Pareto fronts in [15] strongly indicate the presence of first and second order phase transitions. If the information theoretical aspect of the diffusion of messages across the network is considered, those transitions might become thermodynamically relevant. Similarly, the tradeoff of topological robustness when random or targeted nodes are taken away results in a Pareto front [16]. Under the light of our findings, second order transitions are present show up in that study. Also a first order transition exists that vanishes as the average degree of the network changes, suggesting a critical point. We further illustrate our findings with other two examples in section 2.1.

A theory about phase transitions must fit within thermodynamics. For us, this is achieved due to the equivalence between the Pareto front and the Gibbs surface [18], [19], [20], an object known to embody phase transitions in its concavities and non-analyticities. We discuss thermodynamics in section 3.1, not because our theory modifies previous knowledge about it, of course, but because in showing that phase transitions arise in thermodynamics *precisely in the same way* as in MOO, we place our findings for MOO on very solid ground.

2 Theoretical Framework

In this section we expand the loose introduction of MOO above. More details and methods can be found in the exhaustive literature [21], [22], [23], [24], [25]. We assume minimization unless indicated otherwise.

Consider a set X of possible designs $x \in X$. In the example above, $X = \Gamma$ is made of network designs. This will be used again later, along with another example in which $X = A$ stands for all possible languages $a \in A$ derived from a mathematical computational model of human communication [8]. Within X we seek those optimal designs $(x_\pi \in \Pi) \subset X$ that simultaneously minimize a series of *target functions* ($T_f \equiv \{t_1, \dots, t_K\}$). These $t_k \in T_f$ map each design $x \in X$ into *target space* ($T_f(x) = \{t_1(x), \dots, t_K(x)\} \in \mathbb{R}^K$), a morphospace of the system under research.

Pareto dominance is defined in this target space. Take $x, z \in X$. x dominates z (noted $x \prec z$) if $t_k(x) \leq t_k(z)$ for all k and $t_{k'}(x) < t_{k'}(z)$ for at least one k' . This means that x is objectively better than z . If given two designs $(x, y \in X)$ none dominates the other ($x \not\prec y \not\prec x$), we cannot chose one of them without introducing a bias towards some of the target functions. Pareto optimality is solved by putting choices between mutually non-dominated designs on hold.

The *Pareto optimal set* $\Pi \subset X$ is such that every element $z \in X$, $z \notin \Pi$ is dominated by some $x \in \Pi$ while any $x, y \in \Pi$ are mutually non-dominated. The projection $T_f(\Pi)$ conforms a $(D \leq K - 1)$ -dimensional surface in \mathbb{R}^K that embodies the most optimal tradeoff possible between the targets. Moving along

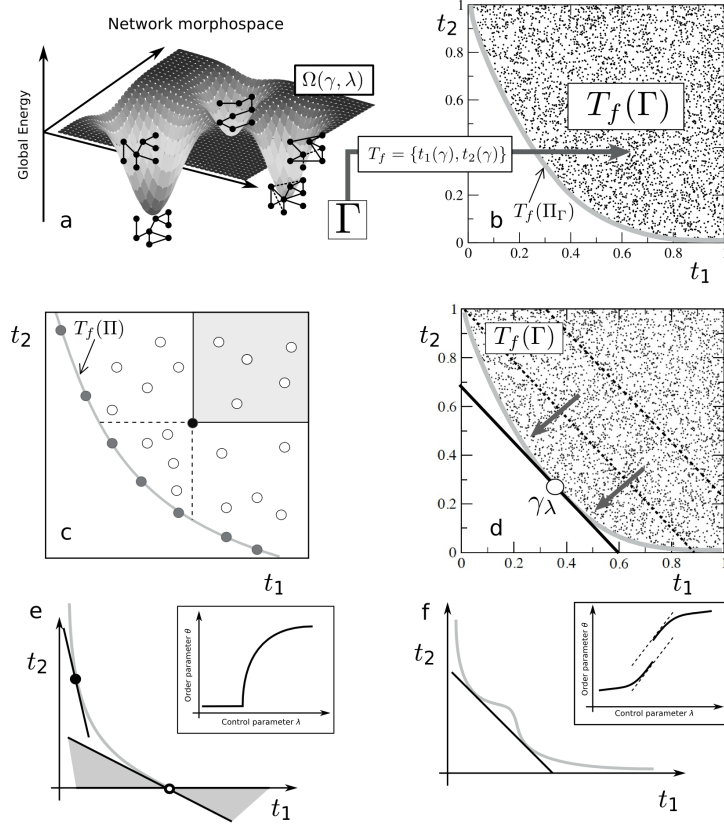


Fig. 1. Phase transitions in Pareto optimal systems. **a** In a design space of complex networks a set of weighted target functions defines a global energy (equation 1) and renders a potential landscape (explored in [18], [13]). **b** Those same targets map the design space into the target space. The set of Pareto optimal designs is mapped onto a boundary of this morphospace: the Pareto front, which represents the most optimal tradeoff between the targets. **c** The concept of dominance is geometrically simple in target space. **d** Energy minimization for fixed λ returns a single point of the Pareto front. Changing λ we visit different solutions. Depending on the shape of the Pareto front, second (e) and first (f) order phase transitions arise as a function of λ .

the front it is impossible to improve all targets at once: an increment in at least one t_k is necessary if we wish to decrease some other $t_{k'}$.

We sketch now the basic situations of our theoretical framework that connects the Pareto front and thermodynamics. We refer the reader to [18] for a more exhaustive discussion.

The simplest SOO problem that includes all MOO targets defines a linear *global energy function*:

$$\Omega(x, \Lambda) = \sum_k \lambda_k t_k(x), \quad (1)$$

where $\Lambda \equiv \{\lambda_k; k = 1, \dots, K\}$ are parameters that bias the optimization towards some of the targets. We say that equation 1 has collapsed the MOO into an SOO. A set Λ with fixed values λ_k defines one single SOO, thus equation 1 (with free λ_k) produces indeed a family of SOOs whose members are parameterized through Λ . We will study: those SOOs, the constraints that the Pareto front imposes to their solutions, and the relationships between different SOOs of the same family. The validity of the results holds for any positive, real set Λ . For convenience, though: i) We take $K = 2$, which simplifies the graphic representations and contains the most relevant cases. ii) We require $\sum_k \lambda_k = 1$ without loss of generality. For $K = 2$ then $\lambda_1 = \lambda$, $\lambda_2 = 1 - \lambda$, and $\Omega = \lambda t_1 + (1 - \lambda)t_2$. iii) We impose $\lambda_k \neq 0 \ \forall k$, thus $\lambda \in (0, 1)$. Comments about fringe cases can be found in [18].

As said above, for given λ_k one fixed SOO problem is posed. Then, equation 1 with fixed Ω defines *equifitness surfaces* noted $\tau_\Lambda(\Omega)$. Each $\tau_\Lambda(\Omega)$ constitutes a $(K - 1)$ -dimensional hyperplane in target space. For $K = 2$ these surfaces become straight lines (figure 1b):

$$\tau_\lambda(\Omega) \equiv \left\{ (t_1, t_2) \mid t_2 = \frac{\Omega}{1 - \lambda} - \frac{\lambda}{1 - \lambda} t_1 \right\}. \quad (2)$$

The slope of $\tau_\lambda(\Omega)$ along each possible direction \hat{t}_k in the target space only depends on λ (here, $dt_2/dt_1 = -\lambda/(1 - \lambda)$). Different $\tau_\lambda(\Omega)$ for fixed λ are parallel to each other. The crossing of $\tau_\lambda(\Omega)$ with each axis is proportional to Ω (from equation 2, the crossings with the horizontal and vertical axes read: $\Omega/(1 - \lambda)$ and Ω/λ). For a given SOO (constant λ), minimizing Ω means finding $\tau_\lambda(\tilde{\Omega})$ with $\tilde{\Omega}$ the lowest value possible such that $\tau_\lambda(\tilde{\Omega})$ still intersects the Pareto front (figure 1d). This is equivalent to *pushing* the equifitness surfaces against the Pareto front as much as possible thus lowering the crossings with the axes.

The SOO optimum $x_\lambda \in \Pi$ lays at the point $T_f(x_\lambda)$ at which $\tau_\lambda(\tilde{\Omega})$ is usually tangent to the front (figure 1d). The exceptions to this rule are the most interesting cases. The solutions to different SOOs (defined by different values of λ) are found in different points along the front. For $\lambda \in (0, 1)$, equifitness surfaces present a slope $-\lambda/(1 - \lambda) = d \in (-\infty, 0)$ (d decreases as λ increases). Consider now differentially small modifications of λ . This allows us to drift infinitesimally

slow through the SOO family. We could expect that solutions between different SOOs will change so gradually as well, but that is not always the case.

The front in figure 1d is convex (with respect to the optimization direction determined by $\lambda \in (0, 1)$). Its slope spans the whole range $d \in (-\infty, 0)$. This guarantees that, as we drift through λ , each different SOO problem has one characteristic solution laying exactly where the equifitness surface is tangent to the front. We can sample the front smoothly, thus anything that we measure on the SOO solutions (i.e. any order parameter) will be a smooth function of λ as well. Convex Pareto fronts whose slope span the whole range $d \in (-\infty, 0)$ do not present any *accident*.

Consider now the case in figure 1e. It represents a convex front whose slope spans $d \in (-\infty, d^* < 0)$. For $\lambda \in (\lambda^* \equiv -d^*/(1 - d^*), 1)$, we can pose different SOOs whose solutions lay at different points of the convex part of the front. Varying λ within this interval renders a smooth sample of SOO solutions. However, for $\lambda \in (0, \lambda^*)$ we can pose different SOOs whose solution lays exactly at the same place, as indicated by the gray fan in figure 1c. If we measure anything about the SOO solutions, that quantity will be constant as a function of λ for $\lambda \in (0, \lambda^*)$ because we will persistently measure a property of the same design. That same property will vary smoothly over $\lambda \in (\lambda^*, 1)$. At λ^* this quantity will be continuous but its derivative will not (figure 1e, inset), as in second order phase transitions. In cases like this we say that the Pareto front ends abruptly at one of its extremes. Second order transitions also happen if the slope of the front spans $d \in (d^* > -\infty, 0)$ (i.e. if the opposite end of the front terminates abruptly) or if $d \in (-\infty, d_-^*) \cup (d_+^*, 0)$ (i.e. the front presents a sharp edge with an ill-defined derivative).

A cavity in the front leads to first order phase transitions. At either side of the cavity in figure 1f we find convex stretches whose points represent different solutions for different SOO problems posed by different $\lambda \in (0, \lambda^*)$ or $\lambda \in (\lambda^*, 1)$. But right at $\lambda = \lambda^*$ (represented by the straight red line of figure 1d) two solutions are SOO optima at the same time. This is a phase coexistence phenomenon characteristic of first order transitions. Pareto optimal solutions laying inside the cavity are bypassed and never get to be SOO optima. If we measure an order parameter of the SOO solutions as a function of λ (figure 1f, inset), we find a gap resulting from the abrupt shift from one convex stretch of the front to the other at $\lambda = \lambda^*$.

2.1 Phase Transitions in Pareto Optimal Designs

As examples, we choose two problems that have recently been treated from an optimization perspective. Take Complex Networks first, which are good models of a series of natural systems such as vascular or nervous circuits [1], [2] that might be constrained by physical costs (available material) while seeking the efficient implementation of biological function (e.g. distribution of nutrients). Some human made structures, such as transportation networks [9], would also benefit from optimal design.

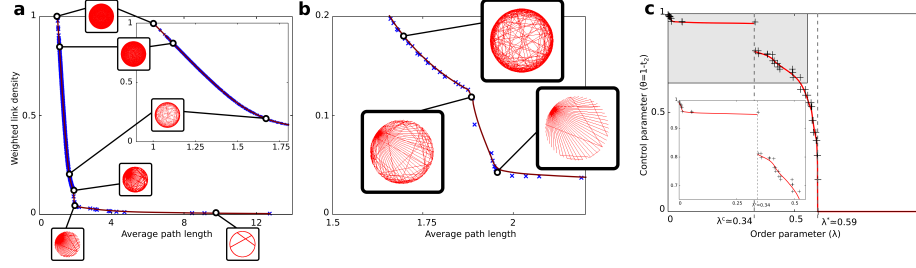


Fig. 2. Pareto optimal networks with nodes spaced over a circle. A genetic algorithm was used to approximate the Pareto front (blue crosses and thick brown curve) of networks that minimize the average path length and the cost of their links. **a** The front implements a tradeoff between the clique and the minimum spanning tree. (Inset) The clique extreme of the front ends abruptly (see [18]) indicating a second order phase transition. **b** A cavity is revealed at the center of the front, which implies a first order phase transition. **c** Both transitions are revealed in the plot of any order parameter. (Second order at $\lambda_2^* \simeq 0.59$, first order at $\lambda_1^* \simeq 0.34$.)

In [13] we consider this problem to a greater extent. We take the cost $\rho(\gamma)$ of network γ as a function of its edges (number or length) and its efficiency is accounted for by the average path length $\langle l \rangle(\gamma)$, a naive proxy for how fast messages can be relayed across the network. These are the targets for minimization ($T_f = \{t_1 \equiv \rho(\gamma), t_2 \equiv \langle l \rangle(\gamma)\}$) that lead to a Pareto front and, depending on its shape, to phase transition and other interesting phenomena. In figure 2a we represent the front for such an MOO along with some Pareto optimal networks. In this example the nodes are spaced over a circle and the cost of each link is proportional to its Euclidean distance. This front ends abruptly (figure 2a, inset) and a cavity is present (figure 2b, see [13] for discussion). This implies, correspondingly, a second and a first order transitions at $\lambda_2^* \simeq 0.59$ and at $\lambda_1^* \simeq 0.34$. These transitions can be noted in the plot of any order parameter (figure 2c).

Our second example explores the evolutionary constraints of human language, an unsettled challenge for the scientific community. The optimization of linguistic structures brings together universal language properties (such as Zipf's law) and the presence of ambiguity, likely as a compromise between language economy and a large ability to talk about the outer world [26]. Such a tension was proposed by Zipf himself [6] and its mathematical formalization [8] leads to an MOO problem that was always treated as an SOO. Accordingly, phase transitions were readily identified but some debate lasted concerning its nature and meaning [27]. In [8], languages $a \in A$ are modeled through a set of signals S and objects R whose associations are encoded in a matrix $a = \{a_{ij}\}$ with $a_{ij} = 1$ if signal $s_i \in S$ names object $r_j \in R$ and 0 otherwise. This binary matrix presents many ones in a row if a signal is polysemous and many ones in a column if several words name the same object – i.e. if there are any synonymous. Every object is recalled equally

often and, if an object has many names, a speaker chooses uniformly among these when necessary. Two quantities are relevant (see [8]): i) one entropy $h_H(a)$ associated to the uncertainty of a message when a hearer has to decode it – i.e. what object it is meant after the speaker has uttered a signal; and ii) another entropy $h_S(a)$ associated to the speaker choosing the right word to name an object among those available. A speaker might be allowed to be vague (as in “it” referring to any object) or she might be requested to be specific (perhaps finding the precise technicism in a scientific context).

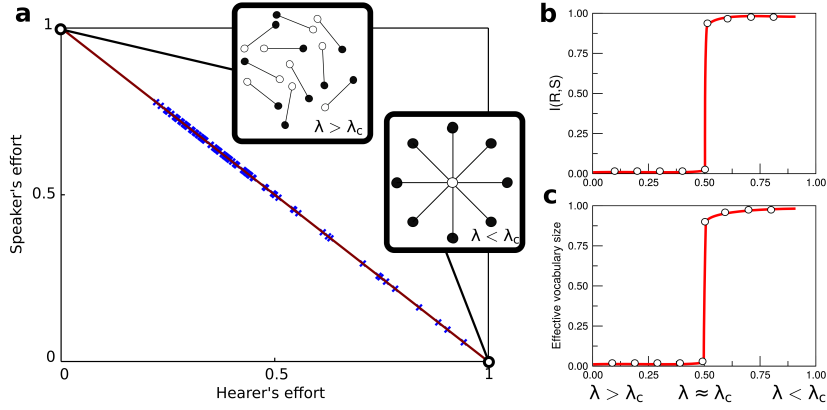


Fig. 3. Least effort languages. **a** Arbitrary Pareto optimal languages (blue crosses) lay on the straight line $t_2 = 1 - t_1$. A straight front is a sign of criticality along a first order phase transition scenario. Either phase represents respectively the best scenario for the speaker ($\lambda < \lambda^c$, where communication is impossible unless through the context) and for the hearer ($\lambda > \lambda^c$, with high memory demands). Only at the critical point is a wide complexity available. Any order parameter (**b**, mutual information between the signals and the external world; **c**, effective vocabulary size) reflects the phase transition.

These two entropies represent the effort made by hearers or speakers when using a language. They act as minimization targets ($T_f = \{t_1 \equiv h_H(a), t_2 \equiv h_S(a)\}$), so that languages $a \in A$ are subjected to a MOO. A subset ($a_\pi \in \Pi$) $\subset A$ of object-signal associations implements the Pareto front, the optimal tradeoff between the efforts of a hearer and a speaker. This front is a straight line in target space (figure 3a). Attending to the theory sketched above this is akin to a first order phase transition (see [18]) with one end of the front being the global optimum for $\lambda < \lambda^c$ and the other extreme of the front being optimum for $\lambda > \lambda^c$. Right at $\lambda = \lambda^c$, a sudden jump happens between these two, very distinct phases. This can be appreciated in any order parameter as a sharp discontinuity (figure 3b, c). The extremes correspond to i) a language that minimizes the speaker's efforts for $\lambda < \lambda^c$ (one single signal names every object, as in the “it” example before, so that the speaker does not need to think

the right association every time) and ii) a language that minimizes the hearer's effort for $\lambda > \lambda^c$, with perfect pairings between signals and objects so that there is not any ambiguity when decoding the messages.

Communication is difficult in both extremes, either because the signals convey little information about the objects ($\lambda < \lambda^c$, figure 3b), or because of the memory needs to browse a vast vocabulary ($\lambda > \lambda^c$, figure 3c). Besides, we know that more complicated structures exist in real languages. These structures can be found right at $\lambda = \lambda^c$. A straight Pareto front is an indication of criticality ([18], [28]). In such cases, exactly at the critical value λ^c , the whole front is also SOO optimal. Note that in usual first order transitions those solutions laying at the cavity are skipped altogether, while here a plethora of them becomes available. In [27] the authors proved that the global SOO minimizers at $\lambda = \lambda^c$ consist of all possible languages without synonyms, hence these must constitute the Pareto front. More importantly, among these possible languages it exists one such that the frequency of the signals obeys Zipf's law, as in natural human languages.

3 Discussion

3.1 Thermodynamics as an MOO-SOO Problem

Thermodynamics is one of the best established branches of physics and dates back to more than two centuries ago. In its modern form – as statistical mechanics – it allows us to make precise predictions about diverse macroscopic physical phenomena. Its applications extend beyond physics, as complex systems are increasingly being investigated through maximum entropy models [29], [30]. In [18] we rewrite thermodynamics as an MOO-SOO problem, not to suggest that our theoretical framework modifies it in any way. Rather the opposite: By checking that our framework reproduces a robust physical theory, we strand our findings in a more solid ground.

Phase transitions in complex systems often raise heated debate: being strict, phase transitions are defined for thermodynamics alone, through partition functions, and involve fluctuations that compel us to take a thermodynamic limit. Little can be done against such epistemological stand. This is yet another reason why we undertake the task of writing thermodynamics as an MOO-SOO. Such a formalization of statistical mechanics reproduces all the results concerning phase transitions *in the exact same way* that transitions arise in other MOO-SOO scenarios. We suggest and support that the phase transition phenomenology arising in other MOO-SOO systems is more than a qualitative similarity.

The argument is not repeated here because of space constraints, but the idea is to show that the *independent, simultaneous* minimization of internal energy and maximization of entropy leads to a Pareto front subjected to the phenomenology found in section 2. In thermodynamics we deal with given physical systems that cannot be modified. We test probabilistic descriptions that tell us how likely it is to find the system in each part of its phase space. We wonder which of these descriptions present a lower internal energy and larger entropy.

Thus our design space X is the set of all possible probabilistic descriptions of the system under research. From that optimization we obtain a Pareto front whose shape (through cavities) and differential geometry (through sharp edges) imply phase transitions if the targets ($t_1 \equiv U$, $t_2 \equiv S$) were collapsed into an SOO problem. But that is precisely what happens in equilibrium thermodynamics through the minimization of the free energy $F = U - TS$ [20]. We identify $\Omega \equiv F$, $\lambda_1 \equiv 1$, and $\lambda_2 \equiv -T$, and the theory exposed above applies with transitions at singular temperature values.

We insist that the optimization operates *upon* probabilistic descriptions of the thermodynamic species – while the shape of the front is determined by the properties of the physical system. It might be interesting to segregate what thermodynamic phenomenology happens because thermodynamic systems *are* probabilistic ensembles (in this regard they are unlike Pareto optimal networks or least effort languages, as much as networks and languages are unlike each other) and what phenomenology arises because of the shape of a Pareto front (that would yield the same phenomenology irrespective of the kind of designs considered – were they networks or languages – as long as the front had the same shape).

This interpretation of statistic-mechanical systems is illustrated with two very simple examples with first and second order transitions and one critical point in [18]. As stated above, this is not to prove new thermodynamic results, but to provide more solid basis for this theory regarding MOO-SOO situations. Indeed, the role of cavities in first order phase transitions dates back to Gibbs [19], [20], whose *Gibbs surface* represented the states of a thermodynamic species. That surface is associated to the microcanonical ensemble [18] and may be concave or convex. Its convex hull is associated to the canonical ensemble (hence to equilibrium at given temperature through free energy minimization), which is always convex. At cavities in the Gibbs surface, the description of both ensembles must differ (as noted by the theory of ensemble inequivalence [31]) and first order transitions occur.

3.2 Closing Remarks

With our recent findings [18] we close a gap between the MOO literature, research on SOO tradeoffs, and statistical mechanics. On the one hand, standard MOO analysis does not take into account phenomena like phase transitions or criticality which, we believe, add up to our knowledge and enrich the description of Pareto optimal designs. On the other hand, analyses of the Pareto front are often qualitative or based on subjective appreciations of its shape. The formalism developed in section 2 allows us to locate quantitatively very relevant details of the systems under research. These features shall persist under transformations of the targets and, if not, the qualitative description would tell us *how* do these phenomena disappear. Furthermore, a solid connection to thermodynamics has been established. We are pretty confident of the immutable, lasting nature of thermodynamics; thus we can guess that, through the Pareto formalism, we have located broad features that unite the description of diverse MOO problems.

A prominent field for MOO application is biology [3], [4], [5]. Thermodynamic-like phenomenology is not discussed in these references, but the stage looks great: Is there a place for true MOO in biology? Against this, natural selection concerns itself with fitness maximization alone. This feels like an exciting MOO-SOO picture, but we cannot guarantee linear global functions as in equation 1. Beyond linearity, new phenomenology might be uncovered.

Finally, an important, though conceptually difficult issue was left aside in [18] and only incidentally dealt with here. How do critical systems look like under an MOO perspective? Can we recover the astounding phenomena of criticality? This will be tackled in future work [28]. Other theoretical aspects of MOO remain open to research.

Acknowledgments This work has been supported by an ERC Advanced Grant, the Botín Foundation, by Banco Santander through its Santander Universities Global Division and by the Santa Fe Institute. We thank CSL members for insightful discussion.

References

1. West, G.B., Brown, J.H., Enquist, B.J.: A general model for the structure and allometry of plant vascular systems. *Nature* 400, 664-667 (1999).
2. Pérez-Escudero, A., de Polavieja, G.G.: Optimally wired subnetwork determines neuroanatomy of *Caenorhabditis elegans*. *Proc. Natl. Acad. Sci. USA* 104(43), 17180-17185 (2007).
3. Shoval, O., Sheftel, H., Shinar, G., Hart, Y., Ramote, O., Mayo, A., Dekel, E., Kavanagh, K., Alon, U.: Evolutionary tradeoffs, Pareto Optimality, and the Geometry of Phenotype Space. *Science* 336, 1157-1160 (2012).
4. Schuetz, R., Zamboni, N., Zampieri, M., Heinemann, M., Sauer, U.: Multidimensional Optimality of Microbial Metabolism. *Science* 336, 601-604 (2012).
5. Szekely, P., Sheftel, H., Mayo, A., Alon, U.: Evolutionary Tradeoffs between Economy and Effectiveness in Biological Homeostasis Systems. *PLoS Comput. Biol.* 9(8), e1003163 (2013).
6. Zipf, G.K.: Human behavior and the principle of least effort. (1949).
7. Maritan, A., Rinaldo, A., Rigon, R., Giacometti, A., Rodríguez-Iturbe, I.: Scaling laws for river networks. *Phys. Rev. E* 53(2), 1510 (1996).
8. R. Ferrer i Cancho, R. V. Solé: Least effort and the origins of scaling in human language. *Proc. Natl. Acad. Sci.* 100(3), 788-791 (2003).
9. Barthelemy, M.: Spatial networks. *Phys. Rep.* 499, 1-101 (2011).
10. Louf, R., Jensen, P., Barthelemy, M.: Emergence of hierarchy in cost-driven growth of spatial networks. *P. Nat. Acad. Sci.* 110(22), 8824-8829 (2013).
11. Dawkins, R.: The blind watchmaker: Why the evidence of evolution reveals a universe without design. WW Norton & Company (1986).
12. Dennett, D.C.: Darwin's dangerous idea. The Sciences (1995).
13. Seoane, L.F., Solé, R.: Phase transitions in Pareto optimal complex networks. <http://arxiv.org/abs/1505.06937> (2015).
14. Avena-Koenigsberger, A., Goñi, J., Solé, R., Sporns, O.: Network morphospace. *J. Royal Soc. Interface* 12(103), 20140881 (2015).

15. Goñi, J., Avena-Koenigsberger, A., de Menizabal, N.V., van den Heuvel M., Betzel, R., Sporns, O.: Exploring the morphospace of communication efficiency in complex networks. *PLoS ONE* 8, e58070 (2013).
16. Priester, C., Schmitt, S., Peixoto, T.P.: Limits and Trade-Offs of Topological Network Robustness. *PLoS ONE* 9(9), e108215 (2014).
17. Otero-Muras, I., Banga, J.R.: Multicriteria global optimization for biocircuit design. *BMC Syst. Biol.* 8, 113 (2014).
18. Seoane, L.F., Solé, R.: A multiobjective optimization approach to statistical mechanics. <http://arxiv.org/abs/1310.6372> (2013).
19. Gibbs, J.W.: A Method of Geometrical Representation of the Thermodynamic Properties of Substances by Means of Surfaces. *Trans. Conn. Acad.* 2, 382-404 (1873).
20. Maxwell, J.C.: *Theory of Heat*. Longmans, Green, and Co., 195-208 (1904).
21. Fonseca, C.M., Fleming, P. J.: An Overview of Evolutionary Algorithms in Multi-objective Optimization. *Evol. Comput.* 3, 1-16 (1995).
22. Dittes, F.M.: Optimization on Rugged Landscapes: A New General Purpose Monte Carlo Approach. *Phys. Rev. Lett.* 76(25), 4651-4655 (1996).
23. Zitzler, E.: *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications*. A dissertation submitted to the Swiss Federal Institute of Technology Zurich for the degree of Doctor of Technical Sciences (1999).
24. Coello, C.A.: Evolutionary Multi-Objective Optimization: A Historical View of the Field. *IEEE Comput. Intell. M.* 1(1), 28-36 (2006).
25. Konak, A., Coit, D.W., Smith, A.E.: Multi-objective optimization using genetic algorithms: A tutorial. *Reliab. Eng. Syst. Safe.* 91(9), 992-1007 (2006).
26. Solé, R., Seoane, L.F.: Ambiguity in language networks. *Linguist. Rev.* 32(1), 5-35 (2014).
27. Prokopenko, M., Ay, N., Obst, O., Polani, D.: Phase transitions in least-effort communications. *J. Stat. Mech.* 2010(11), P11025 (2010).
28. Seoane, L.F., Solé, R.: Systems poised to criticality through Pareto selective forces. In preparation (2015).
29. Harte, J.: *Maximum entropy and ecology: a theory of abundance, distribution, and energetics*. Oxford University Press (2011).
30. Mora, T., Bialek, W.: Are biological systems poised at criticality? *J. Stat. Phys.* 144(2), 268-302 (2011).
31. Touchette, H, Ellis, RS, Turkington, B.: An introduction to the thermodynamic and macrostate levels of nonequivalent ensembles. *Phys A.* 2004(340), 138-146 (2004).